

Date: Mar 05, 2021

EIC Detector R&D Proposal and Progress Report

Project ID: eRD23

Project Name: Streaming readout for EIC detectors

Period Reported: from 6/26/2020 to 2/28/2021

Project Leader: M. Battaglieri and J. C. Bernauer

Contact Person: M. Battaglieri and J. C. Bernauer

Project members

J. Huang, M.L. Purschke

Brookhaven National Laboratory, Upton, NY

S. Ali, V. Berdnikov, T. Horn, M. Muhoza, I. Pegg, R. Trotta

Catholic University of America, Washington DC

M. Battaglieri, M. Bondí, A. Celentano, L. Marsicano, P. Musico

INFN, Genova, Italy

F. Ameli

INFN, Roma La Sapienza, Italy

L. Cappelli, T. Chiarusi, F. Giacomini, C. Pellegrino

INFN, Bologna, Italy

D. K. Hasell, C. Fanelli, I. Frišćić, R. Milner

Massachusetts Institute of Technology, Cambridge, MA

J. C. Bernauer

Stony Brook University, Stony Brook, NY and Riken BNL Research Center, Upton, NY

E. Cline

Stony Brook University, Stony Brook, NY

S. Boyarinov, C. Cuevas, M. Diefenthaler, R. Ent, Y. Furletova, V. Gyurjyan, G. Heyes,

D. Lawrence, B. Raydo

Thomas Jefferson National Accelerator Facility, Newport News, VA

Abstract

The detectors foreseen for the future Electron-Ion Collider will be some of the few major collider detectors to be built from scratch in the 21st century. A truly modern EIC detector design must be complemented with an integrated, 21st century readout scheme that supports the scientific opportunities of the machine, improves time-to-analysis, and maximizes the scientific output. A fully streaming readout (SRO)

design delivers on these promises, however, it can also impose limitations on the sub-detectors and electronics. The streaming readout consortium will research the design space by evaluating and quantifying the parameters for a variety of streaming readout implementations and their implications for sub-detectors by using on-going work on streaming readout, as well as by constructing a few targeted prototypes particularly suited for the EIC environment.

1 Past

What was planned for this period? What was achieved? What was not achieved, why not, and what will be done to correct? How did the COVID-19 pandemic and related closing of labs and facilities affect progress of your project?

The Streaming Readout Workshop VII, organized by BNL, was held in November. Because of the ongoing pandemic, it was purely virtual. The workshop was well attended, with 60 participants registered and over 20 talks. The topics discussed included:

- *Future computing and data rate capabilities.* Modern detector systems, not only in nuclear and particle physics, can already achieve peak rates of TBit/s, requiring to offload the experimental data directly to IT facilities distributed over the campus. In the case of BNL, sPHENIX will be main driver for this network I/O capabilities at SDCC, with the installation of 200-400 GBit/s uplink bandwidth from the sPHENIX counting house to SDCC planned for FY21-22. As current estimates of EIC rates are below those for sPHENIX, on-site network bandwidth will be of no worry. BNL Perimeter connections are likely to be upgraded to 400 Gbit/s in the FY23 time frame, so that off-site data analysis is certainly feasible, and the usage of off-site compute during data taking becomes possible.
- *Streaming readout capable ASICS and platforms from vendors.* In this section, Pacific Microchip, NALU, ALPHACORE and CAEN presented their portfolio and plans for the future. It is clear that the industry moves away from analog buffer designs to more integrated designs with low-power and high-performance ADCs for moderate sample rates up to few 100 MS/s. For higher sample rates, analog buffers are still required. New designs using more advanced buffer handling schemes allow for self-triggering with rate capabilities in the multiple 100 kHz range, making these designs streaming-readout compatible. More and more designs integrate higher-level functions like feature extraction and sub-event building on the FEE chip, lowering the bandwidth requirements of the downstream readout system.
- *Streaming readout at other experiments.* Reports from the CBM, AMBER, sPHENIX, STAR and CLAS12 experiments, all of which have plans for some form of streaming readout, highlight the general trend in the field from traditionally triggered designs to streaming readout. As these experiments plan to run before EIC, they will be important test cases and will provide the community with the opportunity to learn. The experiments run a wide gamut here, from those requiring streaming readout to allow for online track-reconstruction and high-level physics analysis to reduce the amount

of data, to experiments which could run in a more traditional design, but aim for streaming readout to achieve secondary benefits. SPHENIX, for example, will combine the readout of existing, triggered hardware (a cost issue) with streaming readout of other detector parts. The streaming data will be gated by the triggers. However, it is planned to extend the gate time depending on available bandwidth. This allows to parasitically take fully unbiased data, and will allow to access physics not accessible via the triggered setup.

- *Detector support* At the last workshop we tried to involve the various detector R&D groups also supported by this initiative but in the end only eRD14 (PID) and eRD22 (GEM TRD) were able to attend. In addition to these we had a number of status reports and updates about various other experiments running or transitioning to streaming readout systems.

For eRD14, Gary Varner, presented streaming readout considerations for a high performance DIRC, a modular RICH, and a Dual RICH. All of these lend themselves to FPGA readout with streaming capabilities which is being investigated at the University of Hawai and the firm Nalu Scientific. The mRICH is used in Belle and they are currently upgrading both DIRC and mRICH to streaming consistent systems.

For eRD22, Sergey Furlotov, presented the status of a GEM based TRD prototype using a Fadc125 (JLab) with VME readout and a 8 μ s pipeline. With the Fadc125 the GEM TRD can operate as a mini TPC providing 3D track segments. Tests with SAMPA chip found the integration time too long to be useful in TRD applications. They are also investigating machine learning with various analysis architectures: CPU cluster, CPU + GPU, CPU + FPGA, or FPGA only.

Among the detector status reports and upgrades we heard from COMPASS, sPHENIX, GSI, LHCb, and work underway at JLab's Hall B. All of these are developing streaming readout scenarios.

- *SR software* The INDRA-ASTRA project explores the possibility for automated calibrations using AI / ML that would allow a rapid turnaround from data taking to physics results. Abdullah Farhat (ODU) presented a proof of concept for anomaly detection and automatizing calibrations. His studies of anomaly detection are based on detailed detector simulations using his extension of the ADWIN2 algorithm for the higher dimensional data. He reported that the ADWIN2 algorithm has a good balance between a simplicity that leads to computational efficiency and effectiveness in its task. Moreover, its easy extension to the higher-dimensional data of NP experiments, its lack of any prior assumption on the underlying distribution of data samples, and its lack of needing to be trained on any simulated data sets indicates its usefulness. As a next step, the INDRA-ASTRA project will test its concept on a GEM detector prototype in the INDRA facility at Jefferson Lab. Vardan Gyurjyan presented the "ERSAP back-end software for SRO" that is discussed in detail in other sections of this report.
- *Data formats* We discussed the similarities and differences between current data formats employed at the experiments, and how they could be modified for a streaming

readout. Interestingly, most low-level protocols are already inherently streaming read-out compatible or need only minimal modification, as they treat payload data typically as a black box.

One of the SRO-aware software systems currently in production use are the so-called sPHENIX *Event Libraries*, the interface layer between the raw data and user/analysis processes. This layer provides standard access APIs for the data, with an added management component for streaming data. The traditional event paradigm takes on a new meaning with streaming data and becomes a convenient packaging mechanism for a continuous stream, just like, for example, with voice-over-IP applications. The SRO access layer then establishes the management of reassembling the portions of the stream that go together, typically by an embedded clock value.

We continued our strong engagement with the Yellow Report progress, with many of us regularly attending and contributing to the discussions inside the YR DAQ meetings. We help authoring the DAQ R&D and Readout sections of the YR. The YR meetings concluded that a streaming approach is heavily favored for EIC, with no participant speaking out in favor of a traditional readout system. We view this as a strong endorsement for our approach.

Beside the workshop and YR activity, the following group activities, not directly supported by eRD23 funds, were performed. We give here a progress report for completeness:

- *sPHENIX (BNL, SBU)*: The final hardware for computing and networking infrastructure for the sPHENIX DAQ is being installed in the counting-house at RHIC IP8 led by Martin Purschke. The first batch of the hardware includes 96-core DAQ servers, 1.2 PB buffer box servers, and the DAQ network switch, Arista 7280R3. The streaming DAQ for the sPHENIX trackers will be built upon the same hardware in further procurements in FY21. The final production for the sPHENIX FELIX cards continues, which serve as the bridge between streaming front-end and commodity computing. The production will produce 80 BNL-712-version FELIX cards jointly between the sPHENIX and CBM experiments. The max streaming readout bandwidth provided by the FELIX fleet at sPHENIX will be an aggregated 18Tbps.
- *TPEX/DESY test beam (CUA, INFN/JLAB, MIT, SBU)*: Test of a 5×5 array of lead tungstate crystals planned for April 2020 continued to be delayed because of the pandemic and travel restrictions. The DESY test beam facility has and continues to operate with restrictions on occupancy, etc. but we have not made any attempt so far. Hopefully in the summer or fall of 2021 we will be able to resume tests of the prototype calorimeter using three readout schemes in parallel: triggered, INFN Waveboards, and CAEN digitizer.
- *VHDL TDC (SBU)*: A prototype VHDL implementation of a TDC using SERDES IP blocks is in development by an SBU undergraduate. The former student, Marisa Petrusky, left the group. A new student, Kavindra Sahabir, took over this project. Unfortunately, with the pandemic and the steep learning curve of VHDL, not much progress has been made so far.

- *CLAS12 SRO tests (JLab)*: Tests beam were conducted in February and August in Hall-B at Jefferson Lab. They focused on the reaction $e X \rightarrow \pi^0 X$ with $\pi^0 \rightarrow 2\gamma$ produced by the interaction of 10.6 GeV, ~ 100 nA, CEBAF electron beam on $125\mu\text{m}$ lead and 40 cm gaseous deuterium targets. The inclusive π^0 electro-production was chosen because the two decay γ s can be detected by a single detector (the CLAS12 Forward Tagger) reducing the complexity of the experimental set up. The Forward Tagger or FT is part of the CLAS12 detector hosted in Hall-B at Jefferson Lab. It is composed of a lead-tungstate electromagnetic calorimeter (FT-Cal), used to measure the photon energy and position, and a plastic scintillator hodoscope (FT-Hodo), used to distinguish neutrals from charged particles and, in turn, identify gammas. The 332 PbWO crystals of the FT-Cal and the 232 tiles of the FT-Hodo were read out by JLab fADC250 digitizers. The limited number of channels and the combination of two different detectors (calorimeter and plastic scintillators) usually used to trigger the experiment DAQ, represents the ideal bench test for a real on-beam set up. For a quantitative assessment of the streaming read out DAQ chain, some data were also collected in standard triggered mode for later comparison. Off line data analysis is still in progress.
- *fADC250/VTP in streaming mode (JLAB)*: In the Jlab streaming readout (SRO) system physics signals were continuously digitized by the fADC250 flash ADC. The fADC250 is a VME-VXS 16-channel direct-conversion ADC module. These high-speed flash ADCs were developed at JLAB as part of the 12 GeV upgrade. Currently these modules are deployed in many JLAB experiments, providing energy deposition, timing, as well as hit and trigger information. The fADC250 is equipped with an FPGA, that receives 12-bit data-words streaming at 250 MHz from 16 fADC channels in a module. The fADC250 FPGA performs data processing for each fADC channel, computes energy sum of all fADCs, and generates acceptance pulses for each fADC. Each VXS crate houses 16 fADC modules that are managed by the VXS Trigger Processor (VTP) module. The VTP is a VXS switch card module, that was designed to play the leading role in the level1 trigger formation in a traditional triggered DAQ system. The design the VTP contains high speed backplane serial links to each front-end payload module in the crate. Fiber optic serial links provide communication to other crates. In addition, the VTP has more FPGA resources for data processing logic, and a dual-core 1GHz ARM processor, capable of running data processing components, such as event building, trigger and processing diagnostics. These features make the VTP module an ideal candidate for designing a streaming data acquisition system.
- *TriDAS back end software (JLab)*: The Trigger and Data Acquisition System (TriDAS) is software originally designed and implemented for streaming read-out of Astroparticle Physics events. The TriDAS scalable, modular, and flexible design made it also adaptable to the requirements of a beam-based experiment with minimal development effort. TriDAS is made by several software components: the HitManagers (HMs) receive the data streams from a pre-defined number of translators, topologically corresponding to a sector of the detector, subdividing the collected data into a sequence of time-ordered Sector Time Slices (STS); the TriggerCPUs (TCPUs) receive the STS assembled by

all HMs referring to the same TS and apply the event building and the classification/selection algorithms to the data. For the CLAS12 FT application, the Level 1 events (L1) consisted of the detector data within a time window of 200 ns around a hit whose energy exceeded a threshold of ≈ 2 GeV. Level 1 events identified within a TS are then fed to the L2 classification/selection algorithms; a token-based mechanism is at the base of the TriDAS SuperVisor (TSV) load balancing: each TCPU thread owns a token that is given to the TSV on completion of the TS processing; the Event Manager (EM) collects the selected L2 events and then writes them to the so called Post Trigger (PT) file. The TriDAS System Controller (TSC) is the part of the system with which users directly interact configuring and controlling the TriDAS activities. The TriDAS system supports user-level plugins to allow implementation of custom processing algorithms which can be used to implement a software trigger. For Hall-B beam tests, a TriDAS plugin was constructed that implemented the JANA2 framework.

- *JANA2 analysis software (JLab)*: JANA2 is a multi-threaded event processing/analysis framework designed for both offline and streaming applications. User algorithms written within the JANA2 framework were then made available for forming software triggers in the form of JANA2 plugins. The benefit of this is that the full suite of reconstruction algorithms used in the offline reconstruction are available for use as triggers/filters in the streaming system. This includes accessing translation tables, and calibration constants. The software triggering itself was done by using multiple JANA plugins, each implementing their own trigger(s). Each plugin produced one or more *TriggerDecision* objects for each potential “event” identified by the TriDAS system. The decision for each algorithm was in the form of a 16bit integer where a value of zero meant *no-keep* and any non-zero value meant *keep*. If any trigger algorithm indicated a *keep* condition then The TriDAS system was told to keep the event. A unique 16bit ID was assigned to each trigger algorithm (passed in the *TriggerDecision* object). The 16bit ID and 16bit decision for each “event” was given to TriDAS so it could store the decision for each algorithm with each event written out. The JANA2 plugins list was determined by the JANA configuration file. Configuration settings for the individual triggers were also set in this file. The on-demand design of JANA2 specifically supports multi-tiered triggering. This means trigger algorithms can be designed such that more expensive algorithms are only run for events or time slices when a decision cannot be made using the output of less expensive algorithms. The benefit of this is that the compute resource required for the software trigger can be provisioned for the average time needed for a keep/no-keep decision rather than for the most expensive algorithm. An evaluation of the performance of the system during Hall-B tests is in progress.
- *EMCAL (CUA, INFN/JLAB)*: In parallel to the CLAS12 SRO tests in Hall B and in coordination with eRD1 Consortium, a PbWO2 calorimeter prototype has been exposed to a secondary electron/positron beam in Hall-D. The prototype is a matrix of 3x3 crystals made by either PbWO2 or SciGlass. In this test 9 lead-tungstate crystals were coupled to 25um, 6x6 mm² Hamamatsu 13360 SiPMs and connected to the WB2.0, the 250MHz digitizer developed by the INFN Group. With the current set up we obtained a sensitivity of 3.3 pe/MeV deposited. The prototype was calibrated

using cosmic data in the JLab-INDRA Lab and then moved to Hall D. The secondary lepton beam was provided by the Hall D Pair Spectrometer (PS). Each detector arm covers a momentum range of e^+/e^- between 3.0 GeV/c and 6.2 GeV/c. The energy resolution of the pair spectrometer is estimated to be better than 0.6%. The SRO DAQ back end was the same as the one used in Hall-B tests (TriDAS+JANA2). A simplified interface to manage the WB2.0 board during the calibration phase was developed by the INFN group (details were presented at the Streaming RO VI Workshop). During beam tests, signals were recorded from all channels. The off-line data analysis is in progress.

- *Timing module (BNL)*: The global timing module (GTM) to interface sPHENIX with RHIC clock are updated to be based on Xilinx Ultrascale+ FPGAs, in joint development with the control group at BNL NSLS-II. It will handle both clock/synchronization distribution to all subsystems and the global level-1 trigger decision. The prototype clock distribution daughter board interfacing GTM to 24x FELIX cards is being manufactured.
- *SR software (JLAB)* The INDRA-ASTRA project developed a proof of concept for anomaly detection and automatizing calibrations. Their studies of anomaly detection are based on detailed detector simulations using an extension of the ADWIN2 algorithm for the higher dimensional data. The project demonstrated that the ADWIN2 algorithm has a good balance between a simplicity that leads to computational efficiency and effectiveness in its task. Moreover, its easy extension to the higher-dimensional data of NP experiments, its lack of any prior assumption on the underlying distribution of data samples, and its lack of needing to be trained on any simulated data sets indicates its usefulness.
- *Protocol generator (SBU)*: The SBU group started the development of a bit-level protocol documentation and generation tool. Existing protocol/data-format generators like Google's protocol buffers typically allow the user to define the required data fields from a software perspective, while then generating the actual bit-level data layout with minimal or no user control. It is therefore impossible to incorporate given formats, for example defined by a readout chip, into this scheme. In contrast to this approach, the tool under development will allow to define the bit-level data layout, including field descriptions and constraints. From this, the tool then generates publication-level documentation as well as data generators, validators and accessor functions in a series of programming languages. This will facilitate the rapid prototyping of streaming protocols by automatically providing testing tooling and documentation.

How much of your FY21 funding could not be spent due to pandemic related closing of facilities?

Since the pandemic is ongoing, travel is still highly restricted, and we didn't make any use of our funds.

Do you have running costs that are needed even if R&D efforts have paused?

No.

2 Future

What is planned for the next funding cycle and beyond? How, if at all, is this planning different from the original plan? What are critical issues?

The SRO consortium has regular monthly meetings where recent work towards streaming readout is presented and discussed. We are planning the next Streaming Readout Workshop for end of April, hosted virtually by the MIT group. Further, if travel becomes possible again, we plan test beamtimes for streaming readout prototypes at JLAB and at DESY.

Further projects in SRO are listed below. They are not directly funded by eRD23, but some funds might be used for travel in support of these activities:

- The BNL SRO effort will focus on the construction and commissioning of a large-scale streaming tracker system at the sPHENIX experiment. In addition, a new LDRD effort was supported start FY21 in designing ML ASIC aimed at streaming ADC data reduction via neural network-based feature extraction and noise-suppression, which would be applicable for the high sample rate and high background detectors at the EIC (e.g. LGADs).
- CLAS12 SRO test-beams: we are planning to extend tests performed using the Forward Tagger with an improved version of the back end SRO software. The ERSAP/INFN group is actively working to incorporate TriDAS and JANA2 in a single high-performance framework based on CLARA micro-service platform. The JLab SRO system will be tested on-beam using the Forward Tagger to assess the performance. This will be the first step toward the SRO implementation for the full CLAS12 detector. Replacement of front end components not compatible with a SRO operation is in progress (old CAEN TDC are being replaced by VETROCs TDC) and has been planned for the next couple of years. Data collected on-beam will be used to test the full readout chain and the data selection.
- Hall-D SRO calorimeter tests: in collaboration with eRD1 Consortium, a full set of on beam tests using the existing 3x3 crystal prototype are expected. A detailed run plan as been worked out to include different combinations of scintillating material (PbWO2 and SciGlass), readout sensors (SiPMs and PMTs), front end readout (WB2.0 and fADC250), back end framework (ERSAP/INFN CLARA based). Test will be performed at the Hall-D PS facility after testing the systems with cosmic rays in the INDRA Lab.
- Protocol generator and TDC prototype (SBU): The SBU efforts on the protocol generator as well as the FPGA-TDC prototype will continue.
- The VIII Streaming Readout Workshop will be hosted (virtually) by MIT in the last week of April 2021. Format will be similar to previous meetings reporting the ongoing effort within the Consortium and presenting contributions from other Collaborations working on similar SRO projects. A significant fraction of the workshop will be devoted to discuss the SRO solution with the EIC detector efforts to define a strategy to move forward from the R&D phase to final design and prototyping. We expect to have a significant presence of EIC Detectors R&D groups at the workshop.

- The INDRA-ASTRA project will test its concept for anomaly detection and automating calibration on a GEM detector prototype in the INDRA facility at Jefferson Lab. eRD23 is involved in the discussion of the EICUG Software Working Group on data processing and analysis at the EIC. Traditional data models in NP experiments are typically event-oriented and with deep hierarchies reflecting the logical detector arrangement. Considerations of interfacing efficiently to ML and other scientific software tools (such as the SciPy stack), accelerators and modern HPCs, as well as the envisioned streaming readout design for the DAQ, make it important to switch to flat data structures with a column-oriented design to achieve best performance. To this end, for the preparation of the TDR we plan to define common input and output formats for the common software packages to facilitate easy exchange between software components, exchange of software modules, and data preservation. For the completion of the EIC project, in close collaboration with the community involved in the readout design we will define data models suitable for streaming readout and the processing of streamed data in the analysis, both online and offline.
- TPEX test-beams at DESY are still planned for summer or fall 2021 if the travel restrictions from the pandemic are lifted in time. This will include the 5×5 array of lead tungstate crystals with triggered, INFN Waveboard, and CAEN digitizer readout systems in parallel. If the high density glass scintillator being developed at CUA are available we can also exchange some of the PbWO_4 crystals with the glass scintillator.

2.1 Manpower

Include a list of the existing manpower and what approximate fraction each has spent on the project. If students and/or postdocs were funded through the R&D, please state where they were located, what fraction of their time they spend on EIC R&D, and who supervised their work.

All personnel is currently funded by external sources. We report here time spend on SRO related activities, whether directly EIC-related or not.

- SBU/RBRC: Jan C. Bernauer (Assistant Professor) and Ethan Cline (PostDoc) spend about 30% of their time on SRO-related projects. Further, the undergraduate student Kavindra Sahabir is working on the FPGA-based TDC.
- INFN/CUA: the personnel involved in the aforementioned activities at INFN is: Marco Battaglieri (senior staff scientist), Andrea Celentano (staff scientist), Luca Marsicano (post-doc), Mariangela Bondi (post-doc), Simone Vallarino (master thesis student) and Paolo Musico (senior staff engineer) in Genova, F. Ameli (senior staff scientist) in Rome and T. Chiarusi (senior staff scientist), C. Pellegrino (post-doc), F. Giacomini (senior staff scientist) L. Cappelli (post-doc) in Bologna. Each of us has spent approximately 30% of the time on this activity, partially shared on synergistic activities, in particular the BDX experiment at Jefferson Laboratory.
- BNL: J. Huang, M. Purschke will commit 10-20% time developing the SRO system for EIC, which will be supported under BNL LDRD 19-028 and in synergy with on-going

work on sPHENIX SRO tracking system. This work will be supported by an experienced engineering team at BNL including J. Kuczewski, J. Mead, and A. Dellapenna and in collaboration with the ATLAS DAQ team at BNL.

- JLAB: Streaming readout work is performed by staff scientists and engineers from the JLAB DAQ and Electronics groups. D. Lawrence and C. Cuevas are the respective group leaders with B. Raydo, V. Gyurjyan, G. Heyes E. Jastrzembski and J. Gu working 10% of their time on various streaming readout activities. Eric Pooser, a postdoc in Hall-A worked with us in the last year but left for a position in industry. Hall-B and Hall-D DAQ experts S. Boiarinov and A. Somov joined the team. The streaming readout development activities are an ongoing extension of the hardware developed and implemented for the 12GeV experimental programs at JLAB. These activities are collaborative with the SRO consortium groups and complementary work to implement real-time calibration and analysis is also supported by JLAB LDRD-2014.
- MIT: Main effort will be through D. Hasell, R. Milner, I. Frišćić, S. Lee, P. Moran, and B. Johnston working on the DESY test beam. I. Frišćić had a shared role at JLAB and MIT working on streaming readout until the end of 2020. At the beginning of 2021 I. Frišćić returned to Croatia to assume a position there. He will participate in the DESY test beam activities when they resume. C. Fanelli will work together with INFN and JLAB teams to implement Machine Learning and AI-supported high level algorithms within the ERSAP/INFN SRO framework.
- Kentucky: Chris Crawford is the PI of a proposal for the "Kentucky Acquisition and Computational Data Center", which will focus on the development of hardware, firmware and software tools and algorithms for online streaming data analysis using machine learning techniques.

2.2 External Funding

SRO VII was organized by BNL. The INFN Group has been supported by Italian Ministry of Foreign Affairs (MAECI) as Projects of great Relevance within Italy/US Scientific and Technological Cooperation under grant n. MAE0065689 - PGR00799.

Preparation for the DESY test beam activities was supported by MIT's DOE grant DE-FG02-94ER40818.

2.3 Publications

Please provide a list of publications coming out of the R&D effort.

N/A